

DTIC FILE COPY

1

Working Paper No. 377

AD-A224 756

FW/SM: A PROTOTYPE STRUCTURED MODELING ENVIRONMENT

by

Arthur M. Geoffrion

May 1990

DTIC
ELECTE
AUG 6 3 1990
G
D

DISTRIBUTION STATEMENT A	
Approved for public release; Distribution Unlabeled	

WESTERN MANAGEMENT SCIENCE INSTITUTE
University of California, Los Angeles

WESTERN MANAGEMENT SCIENCE INSTITUTE
University of California, Los Angeles

Working Paper No. 377

May 25, 1990

FW/SM: A PROTOTYPE STRUCTURED MODELING ENVIRONMENT

by

Arthur M. Geoffrion

Abstract

A research prototype implementation has been evolving for several years in parallel with the vision, theory, language, and applicative aspects of structured modeling. The objective of this paper is to describe the capabilities of this implementation as it now stands, and to comment on how it contributes toward fulfilling a previously published agenda for the development of a new generation of modeling environments. The intended audience is all modeling system designers and evaluators, including those who do not happen to take a structured modeling approach.

Acknowledgments

Partially supported by the National Science Foundation, the Office of Naval Research, Keton Management Science, Inc., and Shell Development Company. The views contained in this report are those of the author and not of the sponsors.

I am pleased to acknowledge the excellent design and implementation work done by Sergio Maturana, Laurel Neustadter, Yao-Chuan Tsai, and Fernando Vicuña. Their comments on a draft of this paper also were invaluable, as were those of Richard Ramirez and Gordon Wright.

TABLE OF CONTENTS

1. INTRODUCTION

- 1.1 Other Structured Modeling Implementations
- 1.2 Framework III
- 1.3 Background Knowledge

2. OVERVIEW OF FW/SM

- 2.1 Overall Structure
 - FW/SM Programs
 - AUXIL Files
 - Solvers
 - Models
- 2.2 Schema Section
- 2.3 Elemental Detail Section
- 2.4 Reference Section
- 2.5 Activity Section
 - Retrieval
 - Optimization
 - Logical Inference

3. SIGNIFICANCE OF FW/SM

- 3.1 Framework's Desktop Environment
- 3.2 Model Collection
- 3.3 Interfaced Solvers
- 3.4 Model Workspace
- 3.5 The Processes
 - Check General Structure (*SMLCheck*)
 - Check Instantiating Data (*A_RULE* and *UNIQUE*)
 - Check Documentary Comments (*INTERP_CK*)
 - Generate Skeletal Tables (*EDGEN*)
 - Compile Generic Rules into C (*FcEval*) and Execute (*Evaluate*)
 - Generate Reference Documentation
 - Export Tables and Invoke Menu-Driven Query Interface (*XtrieveQuery*)
 - Perform Optimization (*MPS_Interface* and *GENNETFW*)
 - Export Basic Schema Structure and Invoke Arity Prolog (*PrologQuery*)

4. CONCLUSION

- 4.1 Challenge A: Conceptual Framework
- 4.2 Challenge B: Executable Modeling Language
- 4.3 Challenge C: Software Integration
- 4.4 Closing Comments

REFERENCES

Appendix: Automatically Produced Reference Documents

1. Text Documents
2. Tabular Documents
3. Graphical Display

STATEMENT "A" per Charles Mylander
ONR/Code 1111 800 N. Quincy St.
Arlington, VA 22217-5000
TELECON 8/2/90

VG

Accession For	
NTIS	GRA&I
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/ _____	
Availability Codes	
Dist	Avail and/or Special
A-1	



1. INTRODUCTION

The purpose of this paper is to describe and comment on the capabilities of FW/SM, the research prototype implementation of structured modeling (SM) that has been in development at UCLA for about four years. This development has proceeded in parallel with work on requirements for a new generation of modeling environments (Geoffrion <1989c>), on the theory of SM (Geoffrion <1989b>), on a language for SM called SML (for Structured Modeling Language [Geoffrion <1988> and <1990>]), and on various collateral topics.

Although FW/SM was undertaken to support and contribute to this work, this paper is designed to be of use not just to the SM community, but to modeling system designers and evaluators of all persuasions. Some of the features of FW/SM that designers and evaluators may find of particular interest are:

- tightly integrated basic services (viz., business graphics, file management, LAN support, macros, outlining, a rich programming language, spreadsheets, tabular database, telecommunications, and word processing) in addition to specialized modeling capabilities, a feature that promotes FW/SM's usefulness over the whole modeling life cycle rather than at just one part of it
- a menu-driven interface that lies dormant when not in use and thus does not interfere with normal use of the host software package, a feature that illustrates how even sophisticated modeling capabilities can be delivered via popular software
- the use of outlining, boldface, and underlining to relieve users from having to worry about certain technical details of native SML
- unusually comprehensive error-checking -- in fact, exhaustive in a strong sense -- for the formal specification of general model structure, a feature that should help to reduce the time needed to correctly implement, maintain, and enhance models
- an interactive phrase identifier that enables some consistency checking of the natural language comments that SML encourages for documenting formal model specifications, and that also enables some consistency and completeness checking of formal specifications
- automatic generation of several kinds of model reference documentation that facilitate communication, debugging, model maintenance and evolution, and other essential activities
- automatic generation of relational data table designs for model instance data, a feature that facilitates relational database interfaces and achieves a kind of standardization that is essential for efficient implementation of several of FW/SM's capabilities
- automatic evaluation (with warm restart) of indexed families of real- and logical-valued expressions, a feature that facilitates debugging, "What If" studies, and other kinds of model and results analysis
- a graphical browser for the genus graph, a representation of definitional dependencies useful for visualizing the general structure of any model
- a tree-oriented editor for navigation, editing, and display control of the modular outline, an SM construct useful for managing the complexity of complicated models

- complete independence between the general structure of a class of models and instantiating data, a feature that promotes the reuse of each of these, conciseness, ease of communication, dimensional flexibility, and other advantages
- complete independence between models and solvers, a feature that promotes using multiple solvers with a single model, multiple models with a single solver, and conceptual clarity
- a fully automatic interface to a menu-driven commercial database system, a feature that facilitates data/results queries and report generation
- a fully automatic interface for linear and integer programming based on MPS format problem files (no matrix generator need be written)
- a control table interface for generalized network flow optimization (no problem file generator need be written)
- a fully automatic interface to a commercial Prolog system, a feature that facilitates logic-based inferencing with respect to the most fundamental aspects of any model's general structure.

About half of these features are discussed at some length from the perspective of SML in Geoffrion <1990>.

FW/SM has been useful on several levels. At one level, it has enabled experimentation with SM concepts and with SML. At another, it has served as a vehicle for studying implementation issues pertinent to the design of modeling environments. And at still another level, experience with FW/SM has sharpened our vision of what an ideal modeling environment should be like from a user perspective, and of SM's prospects for helping to attain that ideal.

This paper largely avoids operational and implementation details, which are the subject of separate documents.

The organization of the balance of this paper is as follows. The remainder of this introduction surveys the other implementations of SM, reviews the basics of *Framework III* (the system on which FW/SM is built), and indicates what prior knowledge of SM is presumed of the reader. Section 2 and the Appendix present a descriptive overview of FW/SM's functions from the user viewpoint. Section 3 opines on the broader significance of selected aspects of FW/SM. The final section reviews the requirements given in Geoffrion <1989c> for a new generation of modeling environments and explains how FW/SM fulfills most of these within the limitations of a research prototype.

1.1 Other Structured Modeling Implementations

The first implementation of SM ideas, LEXICON (Clemence <1984>), was done in FORTRAN on a mainframe and aimed mainly at supporting optimization for large linear programming problems. The schema and data are entered as batch files. A link to an advanced optimization system was designed but not built.

IIS (for Integrated Information System, Farn <1985>) was built on top of *KnowledgeMan* (Micro Data Base Systems <1985>) on an IBM PC/XT and aimed mainly at demonstrating the feasibility of a hybrid information/analytical modeling system. It includes a menu-driven interface to *LINDO* (Schrage <1987>) for a restricted class of linear programs.

Because LEXICON and IIS were developed early in the evolution of SM and had quite special objectives, they support only a relatively small and nonstandard subset of SML as it exists today.

ASUMMS (for Arizona State University Model Management System, Ramirez <1989>) is a currently active collaborative research effort involving several faculty and graduate students. It adopts SML as its model description language. The architecture includes user, object, relational (*INGRES*), and physical layers. Model/solver/data independence supports the reusability of software and models. Translators map SML to *SAS*, *IFPS*, and *1-2-3*. The hardware environment is a network that spans desktop computers, a VAX 8650, and supercomputers (CRAY X-MP, IBM 3090-500).

Several other specialized or partial implementations of SM exist. These include Hill <1989> (translates schemas from algebra to SML), Jones <1989> (a partially implemented SM specialization of a general graphics-based modeling system), Lenard <1987> (built on top of *KnowledgeMan*, linear programming oriented), O'Dell <1988> (permits graphic SML schema input and translates to relational database form), Park <1987> (ISDOS-based), Worobetz and Wright <1990> (built on top of *ORACLE*, OR/SM has PC-SAS as its primary solver), and Wyant <1988> (translates SML schemas from relational database to graphic form).

With the exception of Worobetz and Wright's OR/SM, the capabilities of each of these implementations have but a modest intersection with the capabilities of FW/SM. Some of the ways in which they differ are worth careful consideration by designers of modeling systems and environments.

In addition, two new UNIX-based implementations are being built at UCLA as part of doctoral dissertations. One, Gokhale <1990>, uses the extensible database system *EXODUS* (Carey and DeWitt <1987>) as its implementation platform. It will exploit novel capabilities that go beyond what traditional database management systems offer, including user defined data types and access methods, enhanced integrity constraints, and version management. The other, Vicuña <1990>, uses attribute grammars to specify formally the complete static semantics of SML, and to generate a syntax-directed editor with incremental type checking and immediate expression evaluation capabilities. It builds on top of the *Synthesizer Generator* (Reps and Teitelbaum <1989>).

1.2 Framework III

FW/SM is built on top of *Framework III*, a well known integrated multi-function program that runs under DOS (Ashton-Tate <1988>). One reason for this choice is that *Framework* makes pervasive use of trees for organizational purposes, as does SM. Another is that it offers many of the basic services needed for a true modeling environment, including outlining, word processing, tabular database, spreadsheets, business graphics, telecommunications, local area networking, and macro recording. It also includes a general-purpose interpreted language called FRED that supports all of these functions. A third reason for choosing *Framework* is that it can transparently run DOS batch files and external programs in other languages, and can move itself temporarily to hard disk or extended memory in order to make more room for loading and executing large programs. This makes it possible to achieve functionalities that would be awkward or inefficient to achieve through *Framework* alone, and to interface *Framework* via files with nearly any stand-alone program.

Prior familiarity with *Framework* would be helpful, but is not a prerequisite for reading this paper because we now briefly summarize the essential concepts needed in what follows. The main ones are: the *Framework* desktop, the tree-oriented editor, and the table editor.

The *Framework* user sees a full-screen "desktop" on which there may be one or more "documents", each of which may be "open" (thus displaying itself in a window that can be resized and dragged) or "closed". There are also ten pull-down menus, and short labels sitting on the desktop that represent the disk drives and closed documents. The user's current selection is shown by a highlight that normally is on or in one of the labels, one of the pull-down menus, or one of the open documents. The general feel is quite Macintosh-like, but with a character-based display. Figure 1 shows an example in which there are two documents -- BUDGET and FEEDMIX -- on the desktop, the second of which is open. Note that there

are five disk drives, plus a "Library" that holds abbreviations, macros, templates, and other tools; clicking on any of these will reveal its contents.

Each document corresponds to one DOS file, and consists of a user-designed tree of "frames". A frame is just a kind of window. The root of the tree corresponds to the document as a whole, each non-leaf node to a "containing frame", and each leaf node to a "word frame" containing text or a "database frame" containing a table whose columns can be named and specified as containing any of six data types. Leaf nodes can also be "spreadsheet frames" or "graph frames", but these kinds are not mentioned further in this paper. The display of a document toggles between an indented outline of its tree called the "outline view", and a "content view" of its tree based on a nested windows representation that reveals the contents of leaf frames. The user can also zoom any leaf frame so that its window fills the screen.

The open FEEDMIX document of Figure 1 is in outline view, from which one can see that its tree of frames has four first-level containing frames labeled "Schema Section", ..., "Activity Section." The children of three of these four frames are hidden, but "Elemental Detail Section" is seen to have six leaf frame children, namely "NUTR", ..., "TOTCOST." It happens that all six are database frames.

Framework's tree-oriented editor, or outliner, for manipulating trees of frames plays a key role. It has facilities to edit tree structure (create, delete, copy, and move nodes and subtrees), to navigate (move toward or away from the root, move to a sibling), and to control display (hide and unhide subtrees or node content).

Framework's table editor, which applies to database frames, also plays a key role in FW/SM. It offers data entry in any of three views (so-called "table", "forms", and "dBASE" views), flexible copy and move operations, row selection ("filtering") for any table via logical expressions and many built-in functions, sorting, column resizing (which can be used to hide columns temporarily), and the ability to compute a variety of column statistics. Word processing facilities also apply to database frames, including the string search, spelling checker, and automatic abbreviation expansion utilities.

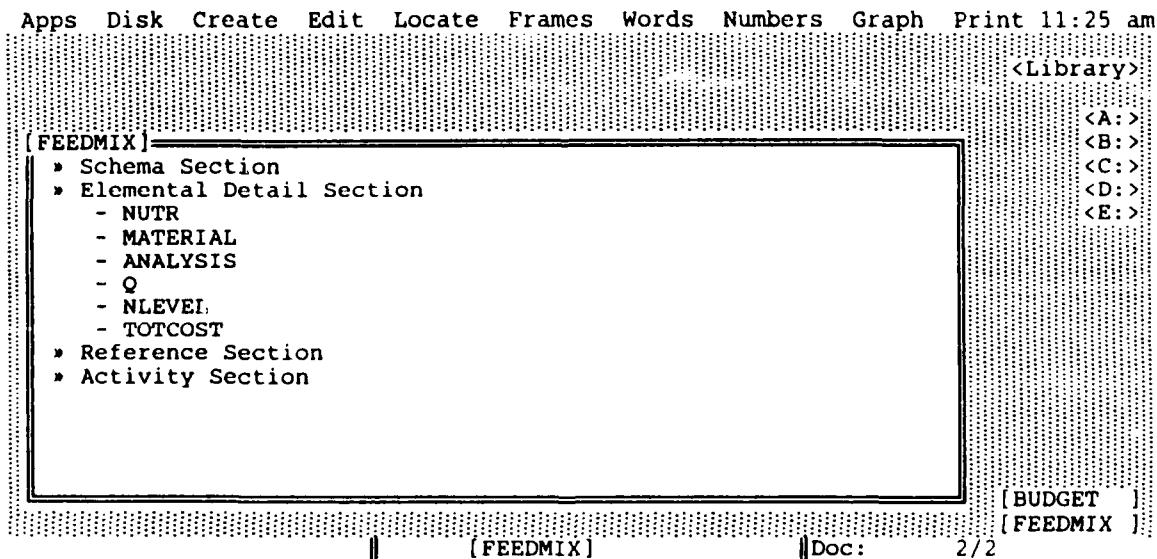


Fig. 1 A Sample Framework Desktop

Framework's ability to communicate with other programs is facilitated by its file import/export capabilities in 9 standard formats. It can export any frame or selected part thereof as a file, and import any file as a frame.

Framework has many other capabilities that need not be described for an understanding of what follows.

Except for those few cases in which FW/SM turns control over to an external program, FW/SM's user interface is *Framework's* user interface. Thus it is no more difficult to use than *Framework*, which has received much critical acclaim.

1.3 Background Knowledge

Full comprehension requires the reader to have prior exposure to SM at the level of Sections 1-3 of Geoffrion <1987>. It is not necessary to have detailed knowledge of SM theory nor of SML beyond what is explained there. It will prove convenient to have a copy of that paper handy while reading this one, as frequent references will be made to its concepts and figures. Because that paper is referenced so often in this one, we abbreviate its reference to "<MS 87>".

The following terms from <MS 87> will be used. All other required concepts are defined as needed. Familiarity with Geoffrion <1989b> would be helpful but is not necessary.

call	genus paragraph
element	genus type (primitive and compound entity, attribute, function, test)
element graph	
elemental detail table	interpretation
elemental detail table stub	key phrase
evaluation	modular outline
externally indexed genus	modular structure
formal part of a genus paragraph	modular tree
generic calling sequence	module paragraph
generic rule	no-forward-reference property of schemas
genus	schema
genus graph	self-indexed genus

2. OVERVIEW OF FW/SM

This section describes the user interface and main capabilities of FW/SM. Most operational and implementation details are omitted, but full details are available in Geoffrion, Maturana, Neustadter, Tsai, and Vicuña <1990>, in a series of internal technical documents, and in a planned companion paper on implementation (Neustadter, Geoffrion, Maturana, Tsai, and Vicuña <1990>). Section 3 appraises the significance of what has been achieved.

FW/SM implements SML as specified in Geoffrion <1988>, and attempts to do so in the spirit of the computer-based modeling environment vision presented in Geoffrion <1989c> (more on that in Section 4). It runs on IBM PC/AT and PS/2 class machines under DOS. We begin with an orientation to the overall file and workspace structure of FW/SM as seen by the user. We then discuss the major components separately.

2.1 Overall Structure

FW/SM offers about 20 special functions beyond those of Framework as menu-selectable processes. The names of these processes appear in *italics* throughout this paper, and will be underlined on first appearance. There is a modest on-line help facility that covers all of the processes.

FW/SM involves four main sets of files:

FW/SM Programs	code that implements the special functions of FW/SM
AUXIL Files	model-specific intermediate files used by certain processes
Solvers	a collection of solvers for model manipulation (presently retrieval, optimization, and inferencing)
Models	model-specific "workspaces" in development or previously developed.

Each is discussed briefly.

FW/SM Programs

FW/SM's programs are written in Framework's FRED language, in C, or in PASCAL. All programs run in a way that is transparent to the user, who sees only a menu of FW/SM processes that appears when a special key is pressed.

AUXIL Files

Certain processes which operate on a schema produce outputs that are not intended for user access, but that are needed by certain other processes. Rather than generating these outputs afresh every time they are needed, FW/SM stores them in a model-specific subdirectory of a subdirectory named AUXIL. This leads to greater client process efficiency so long as the schema does not change.

Solvers

Four solvers have been interfaced to date:

- Xtrieve, an interface and processor for database queries (Novell <1988>)

- GENNET, an optimizer for generalized network flow problems (Brown and McBride <1984>)
- LINDO, an optimizer for linear and integer programming problems (Schrage <1987>)
- Arity Prolog, a language and processor for logic-based inferencing (Arity <1987>).

In each case, there is an unmodified main executable file and associated custom interface programs.

Clearly it is possible to interface other solvers with FW/SM for retrieval, solving simultaneous systems, optimization, inferencing, and other model manipulations.

Models

The store of models will be unique to each installation and perhaps to each user of a modeling environment like FW/SM. At the time of this writing, we maintain a growing library of well over 100 models divided into ten categories: administrative, database, dictionary, economics, engineering, MS/OR, marketing, statistics, test, and miscellaneous. Many are collected in Geoffrion <1989d>. They run the gamut from very simple to moderately complex.

Model files are stored in Framework III's internal format, and can be organized and stored in any convenient DOS subdirectory structure. Each holds a single Model Workspace, the four standard parts of which are the following first-level containing frames, or windows, of the Framework document that appears on the desktop when a file is loaded. See Figure 1.

"Schema Section"	Here the user composes, browses, and maintains a single SML schema.
"Elemental Detail Section"	Here the user loads, browses, and maintains the SML elemental detail tables associated with a single schema.
"Reference Section"	Holds system-generated reference documentation corresponding to the schema and elemental detail; here the user browses and manipulates these exhibits using Framework's editors.
"Activity Section"	Here the user performs model manipulations and related activities throughout the modeling life cycle.

Only a Schema Section is mandatory if any of FW/SM's special processes are to be invoked. The user is free to add other first-level frames to supplement the four just mentioned.

The Model Workspace is the primary face of the system seen by the user. We discuss each of its sections individually.

2.2 Schema Section

An SML schema, recall, is an outline-formatted collection of genus and module paragraphs that describes the general structure of a class of model instances (e.g., Figures 5, 7, and 10 of <MS 87>). Figure 2 shows the schema for a simple feedmix model, to which we shall refer frequently in the balance of this paper. It is an updated version of Figure 5 of <MS 87>.

Given Framework's frame orientation and its versatile outliner for organizing and manipulating frames, it is natural to represent a schema's modular structure as a Framework outline with the Schema

&NUT_DATA NUTRIENT DATA

NUTRi /pe/ There is a list of NUTRIENTS.

MIN (NUTRi) /a/ {NUTR} : Real+ For each NUTRIENT there is a MINIMUM DAILY REQUIREMENT (units per day per animal).

&MATERIALS MATERIALS DATA

MATERIALm /pe/ There is a list of MATERIALS that can be used for feed.

UCOST (MATERIALm) /a/ {MATERIAL} : Real+ Each MATERIAL has a UNIT COST (\$ per pound of material).

ANALYSIS (NUTRi, MATERIALm) /a/ {NUTR} x {MATERIAL} : Real+ For each NUTRIENT-MATERIAL combination, there is an ANALYSIS (units of nutrient per pound of material).

Q (MATERIALm) /a/ {MATERIAL} : Real+ The QUANTITY (pounds per day per animal) of each MATERIAL is to be chosen.

NLEVEL (ANALYSISi, Q) /u/ {NUTR} ; @SUMm (ANALYSISim * Qm) Once the QUANTITIES are chosen, there is a NUTRITION LEVEL (units per day per animal) for each NUTRIENT calculable from the ANALYSIS.

T:NLEVEL (NLEVELi, MINi) /u/ {NUTR} ; NLEVELi >= MINi For each NUTRIENT there is a NUTRITION TEST to determine whether the NUTRITION LEVEL is at least as large as the MINIMUM DAILY REQUIREMENT.

TOTCOST (UCOST, Q) /u/ ; @SUMm (UCOSTm * Qm) There is a TOTAL COST (dollars per day per animal) associated with the chosen QUANTITIES.

Fig. 2 Feedmix Model Schema

Section itself as the containing frame corresponding to the root of the modular structure. Then each module paragraph is also a containing frame, and each genus paragraph is a word frame. With this approach, the outliner's tree-oriented editing facilities are well matched to the tasks of composing and editing modular structure.

What appears in Figure 2 actually is a variant of SML that exploits Framework's word processor and outliner to relieve the user of having to enter reserved words to mark the end of the formal part of each paragraph, to mark the end of each paragraph, to delimit key phrases, and to mark the end of the schema (Geoffrion <1988>). The transition from boldface to non-boldface, the use of underlining, and the use of outlining cause the necessary reserved words to be inserted automatically as a first step in processing the schema.

Two FW/SML processes associated with the Schema Section are of particular interest. The first is SMLCheck, a fully automatic syntax and static semantics checker for the formal part of all SML paragraphs. In addition to checking all syntax, it checks all of the approximately 90 so-called "Schema Properties" given in Geoffrion <1988>. These are context-sensitive static semantic restrictions guaranteeing the integrity of SML schemas. SMLCheck also produces an internal tabular representation of

the schema that is needed by many other processes, and stores it in a model-specific subdirectory of AUXIL.

Whereas *SMLCheck* applies to the formal part of an SML schema, *INTERP_CK* applies to the informal "interpretation" part. It checks whether SML's rules for key phrase syntax and usage have been followed, and it generates an interactive dialog aimed at identifying key phrase variants. To use Figure 2 as an example, each underlined phrase defines a key phrase, and the user must indicate in response to an *INTERP_CK* query that "NUTRIENT" in the MIN paragraph corresponds to "NUTRIENTS" in the NUTR paragraph.

As is true for all FW/SM processes, the user is responsible for invoking *SMLCheck* and *INTERP_CK* by simple menu selection. They produce diagnostics that pinpoint departures from the syntactic and higher order structure prescribed for SML schemas.

In addition, two other FW/SM processes operate on the Schema Section: *FORMAT* sets frame formatting parameters so that printing and the other FW/SM processes will operate correctly, and *COSMET* (for "Cosmetics") trims and sizes all frames to make good use of screen space.

2.3 Elemental Detail Section

SML's elemental detail tables, recall, provide the instantiating data for a single model instance falling within the model class described by the associated schema. Figure 6 of <MS 87> gives sample tables for the schema of Figure 2, and one of these is reproduced near the end of this subsection.

In Framework, it is natural for each elemental detail table to be a database frame. Framework's table-oriented editor is well suited to most keyboard-oriented input and editing tasks arising for elemental detail tables. Framework can import data from external files, but it must be said that more efficient storage structures and links to external files and databases will be needed for many real applications. We comment further on this in Section 4.4.

Two FW/SM processes associated with the Elemental Detail Section are of particular interest, and will be discussed further in Section 3. The first, *EDGEN* (for "Elemental Detail Table Generator"), reads the schema in the Schema Section and creates skeletal elemental detail tables in compliance with the Table Structuring Procedure and Table Naming Conventions of Geffrion <1988>. It is fully automatic unless the user opts to join tables per Step 3 of the aforementioned procedure.

The second process, *FcEval* (for "Function Evaluator"), is a translator that creates a batch-mode evaluator consisting of compiled C code for all function and test genera. This code, which is stored in a model-specific subdirectory of AUXIL, can be executed at any time thereafter by invoking the *Evaluate* process.

FW/SM offers some additional processes applicable to the Elemental Detail Section.

L/E Utilities (L/E for "Loader/Editor") is a collection of utilities that facilitates loading and maintaining elemental detail tables by performing requested relation algebraic operations (e.g., Date <1981>) on their stubs. (Recall from <MS 87> that the "stub" of an elemental detail table comprises the column(s) to the left of the vertical double lines, and that its tuples index the elements of a genus.) Included are utilities for:

- difference (two stubs)
- intersection (two stubs)
- union (two stubs)
- Cartesian product (two or three stubs)

- natural join (two stubs)
- projection (one stub).

In addition, a utility is provided for deleting rows with duplicate stub tuples. These facilities are preparatory to the "smart loader/editor" discussed in Section 4.

Processes are needed to implement the "Table Content Rules" of Geoffrion <1988>, which guarantee the integrity of the elemental detail tables. Two such processes have been implemented to date, both of which pertain to Table Content Rule A. A RULE checks the validity of all identifiers appearing in the stub of elemental detail tables corresponding to externally indexed genera. Such an identifier is "valid" if it is defined in the elemental detail table for its associated self-indexed genus. UNIQUE looks for repeating tuples in the same stub (these are not permissible). About nine more processes will be necessary to fully implement the remaining Table Content Rules.

Finally, MATRIX transforms an elemental detail table with exactly two stub columns and at least one value column into a two-dimensional array in a database frame. For example, MATRIX can transform this elemental detail table associated with the feedmix model's ANALYSIS genus

NUTR MATERIAL			ANALYSIS
=====			
P	std		4
P	add		14
C	std		2
C	add		1

into the array

NUTR	add	std
=====		
P	14	4
C	1	2

or its transpose. Two-dimensional arrays are easier to interpret for some purposes than the linearized version used by elemental detail tables.

2.4 Reference Section

There are many different kinds of reference documentation that one might want to see in connection with a schema, whether it is newly composed or called up from the model files. The same is true, although to a lesser extent, of elemental detail tables. Instead of trying to anticipate and produce all of the documentation that might ever be needed, FW/SM's strategy is to produce several basic machine-resident documents containing the essentials and let the user manipulate them interactively to obtain the specific reports and analyses of interest.

The manipulation can be by Framework's word processor for text documents, by its outliner and word processor for outlines, and by its table editor and word processor for tabular documents. In addition, Framework's rudimentary graph editor can be used for business graphs, although none of FW/SM's special processes produce these.

The reference document names and the processes that produce them are as follows:

Calling Sequence	<i>REFGEN</i>
Genus Graph	<i>NETDRAW</i>

Genus/Module Summary	<i>EDGEN</i>
Identifier Dictionary	<i>ID_DICT</i>
Matrices	
Adjacency Matrix	<i>REFGEN</i>
Reachability Matrix	<i>REFGEN</i>
Adjacency/Reachability Matrix	<i>REFGEN</i>
Modular Outline	<i>REFGEN</i>
Natural Language Summary	<i>REFGEN</i>
Table Structure	<i>EDGEN</i>
Topological Sort	<i>REFGEN</i>

Owing to the considerable space needed to describe these documents and typical manipulations via Framework's editors, we defer further details to the Appendix.

2.5 Activity Section

The Activity Section is the venue of choice for most user interactions with solvers. FW/SM provides one process for retrieval, two for optimization, and one for logical inference. They are discussed in that order.

Retrieval

Three illustrative queries that could arise in connection with feedmix models are:

- (a) List the materials in decreasing order of unit cost.
- (b) List the materials used in quantity greater than 1 pound per day per animal and with unit cost greater than 2 dollars per pound.
- (c) List the materials with above average analysis in those nutrients for which the current mix fails the nutrition test.

The ability to answer such queries is important, over a modeling application's life cycle, in order to gain full advantage of the information that resides in the Elemental Detail Section.

A distinction should be made between queries that involve but a single elemental detail table and those that involve more than one. The first query above involves just one elemental detail table, while the other two involve two tables. Most queries involving a single table can be answered easily with the help of Framework's table editor. But queries involving multiple tables are awkward for Framework, and best handled via FW/SM's interface to the commercial database system Xtrieve (Novell <1988>).

The *XtrieveQuery* process automatically exports the elemental detail tables to Xtrieve in a format that it can understand, invokes Xtrieve with proper initial conditions, and deposits the user in the menu-driven Xtrieve query interface. Upon exiting, the user is returned automatically to FW/SM.

Xtrieve is a quasi-relational system, and so is able to answer most of the usual relational queries, including those that involve multiple tables. The full power of Xtrieve is available to the user, including:

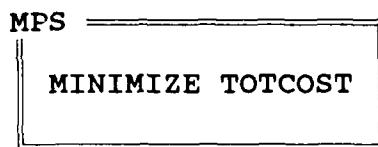
- browsing, sorting, joining, reformatting, and relational selection via complex logical conditions
- the usual aggregate functions for calculating column statistics

- access to a data dictionary summarizing the design of all tables
- the ability to edit the data in any table and to define new tables that may or may not be based on the original elemental detail tables (this does not induce any changes in the Elemental Detail Section, although that would be an easy capability to install)
- a recorder for commonly used keystroke sequences
- various file management and printing facilities, including table export in three different formats.

Optimization

Usually there are many possible optimization problems that a user could pose with respect to a given structured model. The user must, therefore, specify which optimization problem needs to be solved. Different interface styles are possible by which the user may make this specification. One way in which they may differ has to do with how much the user needs to know about the layout of the file(s) required by a target solver. FW/SM has one solver interface that requires the user to know almost nothing, and another that requires the user to know a lot.

The *MPS Interface* process for solving linear and integer programs is fully automatic. The user need only prepare a word frame in the Activity Section titled MPS that contains the keyword MINIMIZE or MAXIMIZE followed by the name of the unindexed function genus corresponding to the objective function. The process assumes that the variables coincide with the variable attribute elements, and, unless otherwise stipulated, that the constraints coincide with the test elements (i.e., they must all evaluate to TRUE). For example, the frame



would specify the standard linear programming problem for the feedmix model.

MPS Interface does the following:

1. Read the schema, elemental detail tables, and MPS frame.
2. Verify that the specified problem is indeed a linear program (possibly with some integer variables) after evaluation of all possibly nonlinear expressions that do not depend on the variable attribute values.
3. Construct an appropriate standard MPS input file.
4. Run LINDO (Schrage <1987>).
5. Import the optimal solution back into Framework with all variable values clearly identified (8 character name length limitations are handled automatically).

This process could be adapted easily to interface any optimizer that accepts standard MPS input files. Also, it can be commanded to stop after step 3, to allow manual control of the solution process from this point on. An extension to nonlinear programming would not be difficult.

The other solver interface, GENNETFW, is for the generalized network flow optimizer GENNET (Brown and McBride <1984>). It requires a schema-specific control table in the Activity Section, a non-procedural device by which the user maps the elemental detail tables into a GENNET input file. This input file organizes the data for arcs in records whose fields correspond to

Tail, Head, Unit Cost, Upper Limit, Lower Limit, Multiplier.

There are conventions for representing the node conservation conditions in records of this same form, including both equality and interval supplies and demands. The user must know these conventions, but they are of no concern here.

An important point is that the control table need never be changed so long as the schema remains the same; inserting, deleting, or updating rows in the elemental detail tables affects neither the content nor the size of the control table. Figure 3 illustrates this table for the classical transportation model appearing in Figures 10 and 11 of <MS 87>. The syntax is straightforward. The column headings are standard for all

TAIL	HEAD	COST	U	L	M
LINK.PLANT	LINK.CUST	LINK.COST			
PLANT.PLANT	PLANT.PLANT		PLANT.SUP		-1
CUST.CUST	CUST.CUST		CUST.DEM	CUST.DEM	

Fig. 3 GENNETFW Control Table for Transportation Model

GENNETFW control tables. Each line generates a group of arcs or node conservation conditions. The first line can be interpreted as follows: generate arcs for all rows in table LINK, with the tail nodes taken from the PLANT column and the head nodes taken from the CUST column; use unit costs from the COST column of table LINK, infinite upper flow limits, zero lower flow limits, and no gains or losses on the flows. The second and third lines generate interval supplies at the plants and equality demands at the customers.

GENNETFW does the following:

1. Read the schema, elemental detail tables, and control table.
2. Construct the GENNET input file specified by the control table.
3. Run GENNET.
4. Import the optimal solution back into Framework with all variable values clearly identified (GENNET's requirement for consecutive integer node names is handled automatically).

Logical Inference

One popular approach to logic-based inference is via the Prolog language, all implementations of which have an integral inference engine. FW/SM's PrologQuery process enables this kind of inference to be accomplished with respect to the basic structural information of an SML schema. It does this by automatically translating this information into Prolog predicates (facts and rules), exporting these along with certain other "stock" predicates to Arity Prolog (Arity <1987>) in a format that it can understand, invoking Arity Prolog with proper initial conditions, and depositing the user in the command-driven Arity Prolog environment. The full power of Arity Prolog is available to the user for answering queries (drawing logical inferences) based on the available schema information. Upon exiting, the user is returned automatically to FW/SM.

The basic structural information available for inferencing is the following:

- 1) all genus and module names,
- 2) the type of each genus,
- 3) the genus graph including, for each arc, which component of the calling genus' generic calling sequence does the calling,
- 4) the modular tree, including the order for each sibling set.

The approach used for translating this information into fact-type predicates suitable for Prolog was devised by Chari <1988>. No additional details of an SML schema, nor any of the information in elemental detail tables, are translated into predicate form; however, it is possible to do so.

Inferencing capabilities are expanded by the "stock" predicates mentioned earlier. These rule-type predicates, also devised by Chari <1988>, capture the relevant schema-related core concepts of Geoffrion <1989b>. It is not so much the final predicates corresponding to these rules that are of use -- for these must be satisfied by any SML schema that passes *SMLCheck* -- but rather the various preparatory predicates introduced to facilitate defining the final predicates.

As an example for those familiar with Prolog, suppose that one wishes to list all primitive entity genera that are called directly or indirectly by a particular genus "gname". To do this, one could define the predicate

```
call_pe(X, Y) :- gcalls_ms(X, Y), agt(Y, pe)
```

and then run it in the form

```
call_pe(gname, Y).
```

Here *agt(Y, pe)* is a fact-type predicate which says that genus *Y* is of primitive entity type, and *gcalls_ms(X, Y)* is a preparatory predicate which is true if genus *X* calls genus *Y* directly or indirectly.

3. SIGNIFICANCE OF FW/SM

The previous section described the overall structure of FW/SM and the function of each of its special processes. This section offers an opinionated discussion of certain features of FW/SM in terms of their broader significance. We are particularly concerned with the implications of FW/SM for designers and evaluators of modeling systems and environments. Thus we incorporate into this discussion many of the non-implementation lessons learned from the FW/SM project.

We take a top-down approach, beginning with Framework's desktop environment and continuing with the model collection, the interfaced solvers, the overall design of Model Workspaces, and some of the processes.

3.1 Framework's Desktop Environment

The decision to build the prototype on top of Framework was among the most important of our implementation decisions. Not only did it have a profound effect on overall system architecture, but it also largely determined the user interface and what basic services would be available to the user. We leave a discussion of the first effect to the planned companion paper on the implementation of FW/SM (Neustadter, et al. <1990>).

Framework's user interface -- reminiscent of Macintosh's and with a strong flavor of direct manipulation (Shneiderman <1987>) -- has proven to be quite easy to use and appropriate to the demands of analytical modeling, at least within the limited range of our experience. This positive assessment is consistent with the current strong trend toward graphical user interfaces, both in general (e.g., PC Magazine <1989>) and in the particular case of analytical modeling software.

Framework's basic services also have proven, in our trials, to be convenient and compatible with the demands of analytical modeling. As advocated in <MS 87> and Geoffrion <1989c>, a modeling environment should support as much of the full modeling life cycle as possible; this in turn requires

... a high degree of software integration, especially with respect to tools and utilities for communication (e.g., business graphics, telecommunications, and word processing or even desktop publishing), for organizing things and ideas (e.g., configuration and version control, database management, file management, outlining, and project management), and for quantitative analysis (e.g., data acquisition, interactive data analysis, graphics, mathematical, spreadsheet, and statistical programs). [Geoffrion <1989c>]

Framework certainly exhibits the requisite degree of software integration, although it does not provide all the services just listed. It does provide outlining, word processing, table editing, file management, macros, and a rich built-in programming language. These have been essential. It also provides spreadsheet modeling, business graphing, and telecommunication facilities, which have proven useful on occasion; and LAN support, which we happen not to have used as yet.

It would be possible to integrate some of the missing services into FW/SM, but it is not necessary to do so for FW/SM users to appreciate the value of well integrated access to tools and utilities like the ones provided. A few workaday examples will underscore this point: (1) Framework's outliner is superb for navigation, editing, and display control of the modular tree that organizes the genus paragraphs in the Schema Section; (2) Framework's table editor is indispensable for working with the tables in the Elemental Detail Section, and for manipulating the documents in the Reference Section; (3) Framework's copy, move, and cut-and-paste facilities find regular use as a way to transfer material among different parts of a Model Workspace, and between Model Workspaces and external files and documents being written with Framework's word processor.

Two other points deserve mention. One is that FW/SM behaves exactly like Framework until the user decides to invoke one of its processes by calling up the special menu. This style, which leaves the user in full control of a familiar commercial software package rather than having to learn a totally unfamiliar modeling system, may turn out to be one of the few avenues leading to much broader acceptance of modeling tools. FW/SM shows that even very sophisticated modeling capabilities can be delivered in this way. Moreover, the modeling environment builder saves the cost of implementing the host package's services which, as we have seen, can be very useful. In a word, FW/SM illustrates a "piggyback" approach to modeling environment design and implementation.

The second point is that modern word processors offer format and font capabilities like outlining, boldface, italics, underlining, subscripts, and superscripts that can be harnessed to improve the readability of model description languages that, for technical reasons having to do with their processing, are limited to a linear string of characters. As explained in Section 2.2, FW/SM exploits some of these capabilities in order to relieve the user of having to enter certain reserved words that serve purely technical purposes. It seems inevitable that format and font options will increasingly be incorporated into future modeling languages.

3.2 Model Collection

DOS subdirectories have proven satisfactory for physically organizing our collection of models. We also use a simple structured model for conceptually organizing this collection. For each Model Workspace, it contains the filename, origin, category, maintainer, domicile, and date of last revision. For maximum usefulness this model should, of course, be maintained automatically by the modeling environment instead of requiring manual updating as at present.

One of the three major design challenges posed in Geoffrion <1989c> is that a modeling environment should be based on a conceptual modeling framework of great generality. As explained there (see also <MS 87> and Geoffrion <1989b>), structured modeling responds to that challenge by formalizing a general class of definitional systems, for which SML furnishes a notation. The diversity of the more than 100 models in our collection confirms the alleged generality. A printed extract from this collection is available in Geoffrion <1989d>.

Some of our models are SML translations of models appearing in the works of other authors. These permit comparisons between SML and FW/SM on the one hand, and alternative modeling styles and systems on the other. A few examples:

- the static Mexican steel industry model from Kendrick, Meeraus, and Alatorre <1983> enables comparison with GAMS (Brooke, Kendrick, and Meeraus <1988>) for economic planning
- the production model from Ellison and Mitra <1982>, Lucas, Mitra, Darby-Dowman <1983>, and Hürlimann and Kohlas <1988> enables comparisons with the UIMP, CAMPS, and LPL systems for linear programming
- the distribution model from Fourer, Gay, and Kernighan <1990> enables comparison with AMPL
- the comprehensive accounting model from LeBlond and Cobb <1985> enables comparison with 1-2-3
- the educational database example from Chapter 27 of Date <1981> enables comparisons with relational, hierarchical, and network database systems
- the shipping industry example from Hammer and McLeod <1981> and the world traveler example from Hull and King <1987> enable comparisons with semantic database systems.

Having a wide variety of models available in one language (rather than in disparate languages) makes it easier for FW/SM users to link together models that have been developed previously (<MS 87>, Geoffrion <1989a>), and to compose new models by starting with previously developed ones much as one might compose a routine business letter by modifying a previous one.

3.3 Interfaced Solvers

FW/SM's automatic interfaces with commercial-quality solvers for retrieval, optimization, and inference demonstrate that models and solvers can be sharply separated, as advocated in <MS 87>.

Moreover, FW/SM shows that a modeling environment can support a wide variety of solvers; e.g., it does not have to cater only to optimizers. This enables synergies that are not possible with just one solver or even one category of solvers. Indeed, we believe that the ability to manipulate a given model in a variety of ways will be one of the characteristics that distinguish today's modeling systems from tomorrow's modeling environments.

3.4 Model Workspace

The division of each Model Workspace into four main sections (windows) has worked out largely as intended. This is particularly true for the Schema Section, Elemental Detail Section, and Reference Section. For example, the cleavage between the Schema Section and the Elemental Detail Section helps to fulfill the objective of general structure/detailed data independence as advocated in <MS 87>.

However, the Activity Section has not been as exclusive a focus for model manipulation activities as originally expected. There seem to be two main reasons for this. First, the *XtrieveQuery* and *PrologQuery* processes transport the user to the user interfaces which these external solvers provide. Second, it is often convenient to create additional first-level frames (sibling windows to the four main ones) for a variety of ad hoc purposes, rather than creating these same frames within the Activity Section containing frame.

This experience suggests that it may be appropriate for modeling systems and environments to accommodate a good deal of flexibility as to how users organize model manipulation activities.

3.5 The Processes

We now comment on the significance of some of FW/SM's processes.

One point that applies to many of the processes is that they have been implemented so as to advance the important objective of general structure/instantiating data independence. Specifically, *SMLCheck*, *INTERP_CK*, *REFGEN*, *NETDRAW*, *EDGEN*, *FcEval*, and *PrologQuery* all work in the absence of any detailed data. Not all extant modeling systems exhibit this sort of independence. For example, GAMS, a leading modeling system for mathematical programming, requires detailed data before a GAMS schema can be processed in any way.

Check General Structure (*SMLCheck*)

With the array of compiler-construction tools now available, especially for important classes of context-free grammars (e.g., Aho, Sethi, and Ullman <1986>), it is now commonplace for the context-free component of a modeling language to be defined formally and used as the basis for the lexical and syntactic analysis parts of their implementation. Thus lexical and syntactic error-checking have become routine, and *SMLCheck* does indeed accomplish this for SML.

But *SMLCheck* does much more than this: it checks the SML Schema Properties. Surprisingly, these properties are *exhaustive* in an important sense. Proposition 2 of Geoffrion <1988> shows that a schema which satisfies all Schema Properties is guaranteed to represent a genuine class of structured models meeting certain desirable qualifications. A Schema Section that satisfies all Schema Properties may not fulfill all of the intentions of the modeler, but at least it is internally consistent and free of essentially all errors that can be checked for in the absence of domain-specific knowledge and meta-knowledge not expressible in SML.

This strong result is possible only because SML was designed to support an explicit semantic framework, that is, a notation-free conceptual formalism for modeling (Geoffrion <1989b>). Such a result is not possible for a language whose semantic framework is implicit, as is the usual case.

A modeling environment like FW/SM whose checks on general model structure are potentially exhaustive in the sense described should make it easier and quicker to achieve correct models than is possible in a modeling environment whose checks are *ad hoc* and not necessarily exhaustive. This same point applies with respect to FW/SM's potential enforcement of exhaustive checks on model instance data, to which we now turn.

Check Instantiating Data (*A_RULE* and *UNIQUE*)

The significance of *A_RULE* and *UNIQUE* is that they represent some of the Table Content Rules, which play the same sort of role with respect to elemental detail tables as Schema Properties play with respect to schemas.

Proposition 1 of Geoffrion <1988> essentially asserts that, if a Schema Section passes *SMLCheck* and a corresponding Elemental Detail Section abides by all Table Content Rules and has 100% fill for all attribute value columns, then the result must be a genuine structured model instance. This model instance may or may not fulfill all of the intentions of the modeler, but -- to repeat a comment made with respect to Proposition 2 -- at least it is internally consistent and free of essentially all errors that can be checked for in the absence of domain-specific knowledge and meta-knowledge not expressible in SML.

FW/SM does not enforce all Table Content Rules at present, although there is no technical reason why it could not be made to do so. However, there is another implementation using a more advanced software development technology for which the enforcement of all Table Content Rules has been designed and largely implemented (Vicuña <1990>). It also implements the enforcement of all Schema Properties. This shows that it is possible for a structured modeling environment to enforce syntactic and semantic restrictions, on both general structure and data, that are complete in the sense of the two propositions mentioned above.

Check Documentary Comments (*INTERP_CK*)

SML imposes certain mild conventions on the interpretation part of genus and module paragraphs (see Appendix 5 of Geoffrion <1988> or Section 2 of Geoffrion <1990>). This encourages modelers to be conceptually precise, and enables *INTERP_CK* to perform several checks on the consistency between the formal and informal parts of a schema. These checks can detect some kinds of inconsistency and incompleteness in the formal part, and inconsistency in the informal part. We believe that schemas which pass these checks are generally higher in quality, more useful, more maintainable, more readable, and better documented than those prepared without documentary comments obeying such conventions.

Our experience has been consistent with this belief. Taking the time to write interpretations for all schema paragraphs usually pays dividends, most obviously in terms of model quality. We have often had the experience of writing a schema quickly without any interpretation at all, only to discover mistakes later when going back and filling in the interpretations. Consequently, we recommend that modeling system

designers consider providing for plain language documentation with a mild but purposeful discipline that is machine-checkable.

Generate Skeletal Tables (*EDGEN*)

The elemental detail table designs created by *EDGEN* are all relations in the sense of relational database theory (e.g., Date <1981>). As mentioned in Geoffrion <1988>, all table names are distinct and all column names are distinct within each table. This sets the stage for the elemental detail tables to be supported by a relational database system. This is an important opportunity, although FW/SM does not truly exploit it since FW/SM's database frames and table editor do not constitute a true "relational database system". The *XtrieveQuery* process, which we discuss later, comes much closer.

For reasons that are widely accepted by the database community, it is desirable for the data tables in a modeling environment to be free of insertion, deletion, and update anomalies arising from functional dependencies (e.g., Date <1981>). Those produced by *EDGEN* essentially are; see Neustadter <1990> for a comprehensive treatment of this result.

A further advantage of FW/SM's tables follows from the fact that their design follows exactly the rules prescribed in Geoffrion <1988>. This standardized table design essentially removes any variability as to how the instantiating data (including calculated values) for a given general structure will be organized. Such standardization is far from ubiquitous among extant modeling systems. About a half dozen of FW/SM's processes depend on it in order to achieve universal applicability.

Compile Generic Rules into C (*FcEval*) and Execute (*Evaluate*)

One of the things that distinguishes analytical modeling from data modeling is the frequent occurrence of mathematical expressions, especially indexed families of expressions. In structured modeling, these families are represented by indexed function and test genera. The generic rule of each such genus gives a typical expression that applies, with appropriate specializing changes, to each element. It is necessary to be able to evaluate all these expressions.

Since the advent of spreadsheets, users are unlikely to be satisfied with any modeling system or environment that fails to provide the immediate feedback of a resident evaluator. Nor should they be, for immediate evaluation is one of the most useful of all debugging tools for both models and results derived by model manipulation (e.g., optimization). It is also an essential facility for answering "What If" questions, for doing many kinds of model analysis and results analysis, and for generating reports. To give just one example, in most applications of optimization, it is essential to be able to recalculate the complete consequences of making manual overrides to the "optimal" solution determined by the solver. See <MS 87> for further motivation.

FcEval automatically produces a suitable resident evaluator. It compiles each generic rule into executable code that can be evaluated for every element of its genus. This code runs whenever the *Evaluate* process is invoked. Note the following points:

1. *FcEval* separates compilation based on general structure from execution based on instantiating data. Thus, it need only be invoked once so long as the Schema Section is not changed in a way that could affect evaluation, and arbitrary edits can be made to the elemental detail tables (insertions, deletions, and updates of all kinds) without having to reinvoke *FcEval*. This "warm restart" capability contributes to the efficiency of *Evaluate* in the usual case where only some of the data change from evaluation to evaluation. Moreover, *FcEval* produces compiled C code rather than interpreted code, which would be less efficient.
2. *Evaluate* puts evaluation results into the elemental detail tables where they are readily accessible, for example, to Framework's table editor and to the *XtrieveQuery* process.

The equivalent of *FcEval* is unusual in today's modeling systems. For example:

- * Database systems can conveniently handle only the simplest kinds of mathematical expressions.
- * Spreadsheet packages usually rely on interpreted rather than compiled code, and require laborious individual entry of expressions in an indexed family unless it happens to be possible to use the copy command to generate the proper family of individual expressions from a single family member.
- * AMPL (Fourer, Gay, and Kernighan <1990>) has no resident evaluator, relying instead on whatever optimizer it is linked to. In any case, non-programmer users must recompile model structure and data *ab initio* if even a single edit is made to the data, and the results of optimizer-executed evaluation are likely to be in a flat file format of limited direct utility.
- * GAMS (Brooke, Kendrick, and Meeraus <1988>) does have a resident evaluator but, like AMPL, non-programmer users must recompile model structure and data *ab initio* if even a single edit is made to the data. Moreover, evaluation results are presented in a flat file format of limited direct utility.

A resident evaluator with the properties mentioned in points 1 and 2 above is an important component of the next generation of modeling systems and environments.

Generate Reference Documentation

FW/SM produces an unusually wide variety of reference documents. The Appendix describes in some detail the documents that the *REFGEN*, *ID_DICT*, *EDGEN*, and *NETDRAW* processes produce. Bear in mind that these documents relate only to the model itself, and are distinct from the kinds of reports produced by solvers (e.g., optimizers produce reports pertaining to an optimal solution).

FW/SM's reference documents are useful for many purposes, including helping modelers see aspects of a model from different viewpoints, debugging model specifications, quickly answering common kinds of queries concerning general structure or details, facilitating communication, and mitigating the ubiquitous problem of inadequate documentation for both modeling and non-modeling professionals.

Most mature modeling systems can produce some sort of reference documentation. By comparison, FW/SM has two main advantages. The first is that, being built on top of an integrated package like Framework, users do not require much computer expertise in order to be able to do some very useful and sometimes quite sophisticated manipulation of the documents, which are all machine-resident as word frames or database frames (except for *NETDRAW*'s display). A number of illustrative manipulations are mentioned in the Appendix. Users don't need to know an operating system and have a stable of utilities and specialized editors, as they probably would with a Unix-based modeling environment like ANALYTICOL (Childs and Meacham <1985>). They just need to know Framework, which was designed for users with limited computer expertise.

The significance of this advantage is not that it encourages future modeling environments to be built using a particular integrated package, but rather that it has enabled us to confirm through personal experience the importance, in the arena of reference documentation, of the kind of software integration called for in the quote given in Section 3.1.

The second advantage of FW/SM is that it is based on a conceptual modeling framework that gives a central role to capturing definitional dependencies when specifying the general structure of a model. Thus it is possible for FW/SM to produce documents that cannot be produced without this kind of information, namely (at present): the Calling Sequence report, the Topological Sort, the Adjacency and Reachability Matrices, and (a version of) the Genus Graph. We remark that definitional dependency information is valuable not only to enable important kinds of reference documentation, but also for such uses

as a) model maintenance and evolution, the efficiency of which depends importantly on knowing all implications of any given change, and b) integrated modeling (see Section 3.2 of <MS 87>).

Export Tables and Invoke Menu-Driven Query Interface (*XtrieveQuery*)

The *XtrieveQuery* process demonstrates the possibility that similar automatic interfaces to relational DBMSs may be achievable for any modeling environment that separates its specification of general structure from its specification of instantiating data, and that organizes its data in relations with predictable structure. The desirability of such interfaces was argued in <MS 87>.

A DBMS interface gives access to a query interface and query processor for efficiently answering a wide variety of questions with respect to a model's data, optimal variable values (if an optimizer has been invoked), and other calculated values. The lack of such query capabilities in most extant modeling systems constitutes one of the greatest anachronisms of contemporary modeling practice.

Perform Optimization (*MPS_Interface* and *GENNETFW*)

The *MPS_Interface* and *GENNETFW* processes demonstrate that optimization-oriented modeling systems like those for the AMPL (Fourer, Gay, and Kernighan <1990>), GAMS (Brooke, Kendrick, and Meeraus <1988>), and LINGO (Cunningham and Schrage <1989>) languages may have the potential to evolve into modeling environments that support much more than just the optimization phase of a modeling project. This would bring about a qualitative improvement in their usefulness because they could then support more of the modeling life cycle.

We note in passing that Geoffrion <1989e> contains a comparison of these three languages and SML from the viewpoint of indexing structures, which is one of the most fundamental features possessed by most modeling languages for optimization.

FW/SM's two optimization interfaces are not designed for large-scale problem solving, but -- as pointed out in Section 2.5 -- they do enable experimentation with two different design positions about how much the user is required to know. It is not clear at this time which style of interface is best in which circumstance. There are advantages to requiring the user to understand both the model and the desired solver well enough to be able to fill out a simple control table linking the two. On the other hand, there are advantages to minimizing the work required of the user. This is an area for future study.

Whichever interface style is adopted, no computer programming skill should be required of the user, and it should be easy to switch optimizers if more than one is available that applies to the model at hand. These are properties that surely will be found in all successful modeling environments of tomorrow.

Export Basic Schema Structure and Invoke Arity Prolog (*PrologQuery*)

Support for the process of model formulation is one of the weak links of today's modeling systems. This process requires a great deal of ad hoc reasoning about partial or interim model structures, and much additional ad hoc reasoning takes place with respect to the general structure of models as they are instantiated, communicated, solved, and used in various other ways throughout the modeling life cycle. Thus there is a clear need for the kind of logical inference capabilities that *PrologQuery* makes available for ad hoc reasoning about the most fundamental aspects of general model structure. Such capabilities are unusual among contemporary modeling systems, except for those based on logic modeling.

FW/SM's simultaneous support for logical inference together with solvers for retrieval and optimization presages the advent of modeling environments that offer inferencing capability in addition to the more usual solution capabilities around which current modeling systems are built.

4. CONCLUSION

The main task of this section is to appraise the degree to which FW/SM fulfills the vision presented in Geoffrion <1989c> for a new generation of modeling environments. That article, it should be noted, only mentions structured modeling in passing.

The article argues that a modeling environment needs five properties in order to achieve the primary goals of higher quality, higher productivity, and greater popularity for model-based work. It should:

- 1) nurture the entire modeling life cycle, not just part of it;
- 2) be hospitable to decision and policy-makers, not just to modeling professionals;
- 3) facilitate the maintenance and ongoing evolution of the models and systems built within it;
- 4) enable all of its inhabitants to "speak" the same paradigm-neutral language for model definition;
- 5) facilitate good management of key resources, namely data, models, solvers, and results derived from these.

The article goes on to detail the design implications of these properties, and therefrom to infer three major design challenges:

- A) the need for a general conceptual framework for modeling;
- B) the need for an executable model description language based on this framework;
- C) effective software integration covering (i) linkable libraries of data sources, models, solvers, and derived results, (ii) programs that achieve the necessary kinds of "executability", and (iii) the tools and utilities needed for total life cycle support of modeling work.

We discuss each of these challenges.

4.1 Challenge A: Conceptual Framework

<MS 87> and Geoffrion <1989b> put forth structured modeling as an answer to Challenge A. This framework is demonstrably general in applicability, rigorous, paradigm-neutral and yet compatible with most modeling paradigms in MS/OR, and suitable for use as a foundation for meeting Challenge B. See, for example, the discussion of Section 3.2 of the present paper. These are four of the six requirements that Geoffrion <1989c> lists for a framework to meet Challenge A. A fifth is that it should be "understandable and natural for the main players at each stage of the modeling life cycle". It is too soon in the history of structured modeling to reach a definitive judgement on this requirement, but the appearance of FW/SM and the new SML tutorial set the stage for a proper evaluation.

The final requirement is that the conceptual modeling framework should be consistent with "good" modeling style. We believe that structured modeling meets this admittedly subjective requirement, and arguments to this effect have been made elsewhere. Some of the reasons: it separates general structure from instantiating data, separates models from problems from solvers, exploits repetitive structure (that is the main task of genera), organizes models hierarchically in modules, and documents interdependencies.

FW/SM's role with respect to Challenge A is to offer the opportunity for hands-on experimentation with one possible realization of the structured modeling framework, thereby enabling individuals to judge for themselves this framework's potential with respect to the six requirements.

4.2 Challenge B: Executable Modeling Language

Geoffrion <1988> puts forth SML as an answer to Challenge B and studies its theoretical adequacy. We mentioned two of the results of that study in Section 3.5 of the present paper.

Geoffrion <1989c> points out that five of the six requirements relating to Challenge A also apply to any executable modeling language aspiring to support the chosen modeling framework. Because SML was designed specifically to support the structured modeling framework, it was possible to design SML so that it satisfies these requirements to about the same degree as the structured modeling framework satisfies them in the context of Challenge A.

FW/SM proves that SML is "executable", but does it provide the necessary functionality? Geoffrion <1989c> lists four desirable kinds of functionality as error-trapping, automatic documentation, solver interface setup, and a smart loader/editor for detailed data. We discuss each in turn.

Error-Trapping The theory behind the *SMLCheck* process shows that syntactic and semantic error-trapping can be complete in an important sense at the level of model classes, as explained in Section 3.5. A similar statement holds relative to error-trapping at the level of detailed data, although FW/SM is as yet quite incomplete in this regard. Vicuña <1990> shows that complete error-trapping is possible at both levels.

Automatic Documentation Thanks to *REFGEN* and a few other processes, FW/SM can automatically produce a variety of useful kinds of model documentation.

Solver Interface Setup FW/SM's ability to automatically set up easy, no-programming interfaces to external solvers for optimization, query, and logical inference is demonstrated by the *GENNETFW*, *MPS_Interface*, *XtrieveQuery*, and *PrologQuery* processes.

Smart Loader/Editor A smart loader/editor for detailed data should be set up automatically, be customized to the general model structure at hand, and should offer: (a) suitable data structures together with a passive interface for entry and editing, and (b) an active interface for data entry and editing. The active interface should be able to: (b1) automatically populate indexing structures that are fully defined in terms of other indexing structures (e.g., enter a Cartesian product where the general structure calls for one), (b2) diagnose what operations are needed to achieve a state of internal data consistency, and (b3) carry out, or at least assist in performing, the needed operations. FW/SM offers (a) via *EDGEN*. It does not yet offer (b). The technology of *FcEval* can be adapted to achieve (b1). Function (b2) is possible because SML aims to make all definitional dependencies explicit via the generic calling sequence, and because the Table Content Rules of Geoffrion <1988> constitute a comprehensive characterization of internal data consistency (as explained in Section 3.5 of this paper). Function (b3) can build on functions (b1) and (b2). Detailed design is under way.

In summary, FW/SM's main role with respect to Challenge B is to demonstrate that most (perhaps all) of the necessary executability can be achieved, and to enable SML's worth to be judged on the basis of hands-on experience.

4.3 Challenge C: Software Integration

FW/SM illustrates one possible approach to achieving the level of software integration called for by Challenge C. Namely, add new modeling capabilities to a popular software package that already provides a highly integrated set of services needed by any modeling environment. The first three sections of this paper show that FW/SM adequately meets Challenge C at the level of a research prototype.

The forthcoming companion paper on FW/SM's implementation (Neustadter et al. <1990>) will discuss FW/SM's approach to software integration from a programming viewpoint. More generally, it would appear that the technical problems posed by Challenge C are rapidly being mitigated by intense commercial activity in the software development community aimed at database connectivity, graphical user interfaces, interoperability, multitasking, networking, and more advanced application development tools and operating systems.

4.4 Closing Comments

It should go without saying that the modeling environment ideals presented in Geoffrion <1989c> are not immutable, but rather serve as a point of departure. We agree strongly with the viewpoint expressed in Barstow et al. <1984> in the context of programming (not modeling) environments. The following quote from the editors' closing chapter is equally applicable in the present context:

The process of developing programming environments is highly experimental. We have no hope of abstractly defining the "ideal" programming environment, which can then be implemented. Our best strategy, in fact our only viable one, is to continue to develop environments, experiment with them to identify those aspects which are either helpful or useless, and to incorporate these lessons into the next experiment.

We have approached FW/SM in this spirit. We have used it to confirm theoretical results (e.g., various results about the internal consistency of SML), to probe conjectures (e.g., concerning the generality of SML and the usefulness of modular structure for managing complexity), to test the feasibility of design options (e.g., three major options for designing an evaluator for function and test genera), and to demonstrate structured modeling ideas to outsiders in quest of their feedback. We have also used it to experimentally guide the evolution of our ideas about modeling environments. In fact, SML itself and the vision presented in Geoffrion <1989c> were influenced to a considerable degree by experience with an evolving FW/SM.

Our plans for FW/SM's future include extending its capabilities, using it to study the learnability and effectiveness of SML for various classes of users, and studying the applicability of FW/SM and the ideas it incorporates to practical modeling work and to the evolutionary improvement of other modeling systems. Our experience with FW/SM, together with the experience of others who are presently exploring alternative design concepts and implementation technologies, will guide our future efforts relative to FW/SM and new research prototypes.

To mention just one specific project for extending FW/SM's capabilities, we are planning to provide an option whereby data can be stored in an external database system rather than in Framework. All but two minor processes can adapt easily to this change, which will make it possible to work efficiently with models having one or two orders of magnitude more data than would be practical for Framework to handle.

In closing, it should be said that evolving research prototypes of this complexity are very difficult to sustain in a university setting. The demands for professional skills, financial resources, and sheer patience are severe. Nevertheless, it is this writer's firm belief that research on modeling environments requires experimental developments of this sort if it is to achieve its full potential for success.

REFERENCES

AHO, A.V., R. SETHI and J. ULLMAN <1986>. *Compilers: Principles, Techniques, and Tools*, Addison-Wesley, Reading, MA.

ARITY <1987>. "Arity Prolog Version 5.0," Arity Corporation, 30 Domino Drive, Concord, MA 01742.

ASHTON-TATE <1988>. *Framework III*, Ashton-Tate, 20101 Hamilton Ave., Torrance, CA 90502.

BARSTOW, D.R., H.E. SHROBE and E. SANDEWALL (eds.) <1984>. *Interactive Programming Environments*, McGraw-Hill, New York.

BROOKE, A., D. KENDRICK and A. MEERAUS <1988>. *GAMS: A User's Guide*, The Scientific Press, Redwood City, CA.

BROWN, G.G. and R.D. McBRIDE <1984>. "Solving Generalized Networks," *Management Science*, 30:12 (December), pp. 1497-1523.

CAREY, M.J. and D.J. DeWITT <1987>. "An Overview of the EXODUS Project," *Database Engineering*, 10:2 (June), pp. 47-54.

CHARI, S. <1988>. *Knowledge Representation Using Structured Modeling*, Ph.D. Thesis, Anderson Graduate School of Management, UCLA, 225 pages.

CHILDS, C. and C.R. MEACHAM <1985>. "ANALYTICOL - An Analytical Computing Environment," *AT&T Technical J.*, 64:9 (November), pp. 1995-2007.

CLEMENCE, Jr., R.D. <1984>. *LEXICON: A Structured Modeling System for Optimization*, Master's Thesis, Naval Postgraduate School, Monterey, CA, June.

CUNNINGHAM, K. and L. SCHRAGE <1989>. "The LINGO Modeling Language," University of Chicago, 89 pages.

DATE, C.J. <1981>. *An Introduction to Database Systems*, Volume 1, Third Edition, Addison-Wesley, Reading, MA.

ELLISON, E.F.D. and G. MITRA <1982>. "UIMP: User Interface for Mathematical Programming," *ACM Trans. Math. Software*, 8:3 (September), pp. 229-255.

FARN, C.K. <1985>. *An Integrated Information System Architecture Based on Structured Modeling*, Ph.D. Thesis, Graduate School of Management, UCLA.

FOURER, R., D.M. GAY and B.W. KERNIGHAN <1990>. "A Mathematical Programming Language," *Management Science*, 36:5 (May), pp. 519-554.

GEOFFRION, A.M. <1987> (MS 87). "An Introduction to Structured Modeling," *Management Science*, 33:5 (May), pp. 547-588.

GEOFFRION, A.M. <1988>. "SML: A Model Definition Language for Structured Modeling," Working Paper 360, Western Management Science Institute, UCLA, 129 pages, May. Revised November, 1989.

GEOFFRION, A.M. <1989a>. "Reusing Structured Models via Model Integration," *Proceedings of the Twenty-Second Annual Hawaii International Conference on System Sciences* (held in Kailua-Kona, January 1-3), IEEE Computer Society Press, Washington, pp. 601-611.

GEOFFRION, A.M. <1989b>. "The Formal Aspects of Structured Modeling," *Operations Research*, 37:1 (January-February), pp. 30-51.

GEOFFRION, A.M. <1989c>. "Computer-Based Modeling Environments," *European Journal of Operational Research*, 41:1 (July), pp. 33-43.

GEOFFRION, A.M. <1989d>. "A Library of Structured Models," Informal Note, Anderson Graduate School of Management, UCLA, 200 pages, November.

GEOFFRION, A.M. <1989e>. "Indexing in Modeling Languages for Mathematical Programming," Working Paper 371, Western Management Science Institute, UCLA, 180 pages, November. Portions appeared as "A Taxonomy of Indexing Structures for Mathematical Programming Modeling Languages," in *Proceedings of the Twenty-Third Annual Hawaii International Conference on System Sciences* (held in Kailua-Kona, January 2-5, 1990), Volume III, IEEE Computer Society Press, Washington, pp. 463-473.

GEOFFRION, A.M. <1990>. "The SML Language for Structured Modeling," Working Paper 378, Western Management Science Institute, UCLA, 155 pages, June.

GEOFFRION, A., S. MATORANA, L. NEUSTADTER, Y. TSAI and F. VICUÑA <1990>. "User Documentation for FW/SM Release X90-02," Anderson Graduate School of Management, UCLA, 75 pages, February.

GOKHALE, S. <1990>. "Extensible Databases as a Platform for Modeling," forthcoming Ph.D. dissertation, Anderson Graduate School of Management, UCLA.

HAMMER, M. and D. McLEOD <1981>. "Database Description with SDM: A Semantic Database Model," *ACM Trans. Database Systems*, 6:3 (September), pp. 351-386.

HILL, D.S. <1989>. "A Prototype for Converting Linear Programming Models to Structured Modeling Graphs," Master's Thesis in Information Systems, Naval Postgraduate School, 137 pages, March.

HULL, R. and R. KING <1987>. "Semantic Database Modeling: Survey, Applications, and Research Issues," *ACM Comp. Surveys*, 19:3 (September), pp. 201-260.

HÜRLIMANN, T. and J. KOHLAS <1988>. "LPL: A Structured Language for Linear Programming Modeling," *OR Spektrum*, 10, pp. 55-63.

JAVELIN <1986>. *Javelin*, Javelin Software Corporation, One Kendall Square, Building 200, Cambridge, MA 02139.

JONES, C.V. <1989>. "Graph-Based Modeling and Structured Modeling," Working Paper 89-03-05, The Wharton School, Univ. of Pennsylvania, 37 pages, June.

KENDRICK, D.A., A. MEERAUS and J. ALATORRE <1983>. *The Planning of Investment Programs in the Steel Industry*, The Johns Hopkins University Press, Baltimore, MD.

LeBLOND, G.T. and D.F. COBB <1985>. *Using 1-2-3*, Second Edition, Que Corporation, Indianapolis.

LENARD, M. <1987>. "An Implementation of Structured Modeling Using a Relational Database Management System," paper presented at the Second Annual Conference on Integrated Modeling Systems, University of Texas, Austin, 20 pages, October.

LUCAS, C., G. MITRA and K. DARBY-DOWMAN <1983>. "Modeling of Mathematical Programs: An Analysis of Strategy and an Outline Description of a Computer Assisted System," Report TR/09/83, Department of Mathematics and Statistics, Brunel University, U.K., October.

MICRO DATA BASE SYSTEMS <1985>. *KnowledgeMan*, Micro Data Base Systems, Inc., P.O. Box 248, Lafayette, IN 47902.

MS 87. See Geoffrion <1987>.

NEUSTADTER, L. <1990>. "On the Structure of Data in SML Models," Research Paper, Anderson Graduate School of Management, UCLA, 62 pages, March.

NEUSTADTER, L., A. GEOFFRION, S. MATURANA, Y. TSAI and F. VICUÑA <1990>. "The Design and Implementation of a Prototype Structured Modeling Environment," Working Paper 380, Western Management Science Institute, UCLA, in preparation.

NOVELL <1988>. *Xtrieve: Interactive Query Manual*, Novell Development Products Division, 6034 W. Courtyard, Suite 220, Austin, TX.

O'DELL, D.D. <1988>. "The Design and Implementation of a Visual User Interface for a Structured Model Management System," Master's Thesis in Information Systems, Naval Postgraduate School, 170 pages, March.

PARK, J.Y. <1987>. "A Computer-Aided Tool for Decision Support Model Building Based on Structured Modeling," Master of Engineering Thesis, Korea Advanced Institute of Science and Technology, 57 pages plus appendices, December.

PC MAGAZINE <1989>. "Graphical User Interfaces" (Cover Story), *PC Magazine*, 8:15 (Sept. 12), pp. 97-157.

RAMIREZ, R. <1989>. "The ASUMMS Project: An Overview," ASUMMS Informal Note No. 1, Dept. of Decision and Information Systems, Arizona State University, 14 pages, August.

REPS, T.W. and T. TEITELBAUM <1989>. *The Synthesizer Generator*, Springer-Verlag, New York.

SCHRAGE, L. <1987>. *User's Manual for Linear, Integer, and Quadratic Programming with LINDO*, Scientific Press, Redwood City, CA.

SHNEIDERMAN, B. <1987>. *Designing the User Interface*, Addison-Wesley, Reading, MA.

VICUÑA, F. <1990>. "Semantic Formalization in Mathematical Modeling Languages," Ph.D. dissertation, Computer Science Department, UCLA, June.

WOROBETZ, N.D. and G. P. WRIGHT <1990>. "A Computer-Based Management System for Statistical Models," paper presented at the Joint National TIMS/ORSA Meeting, Las Vegas, May 7-9. Written paper in preparation.

WYANT, M.A. <1988>. "Design and Implementation of a Prototype Graphical User Interface for a Model Management System," Master's Thesis in Information Systems, Naval Postgraduate School, 98 pages, March.

Appendix

AUTOMATICALLY PRODUCED REFERENCE DOCUMENTS

This Appendix describes the basic documents that three of FW/SM's processes produce on request and place in the Reference Section. It also gives examples of typical manipulations via Framework's editors. Most of the documents are produced by REFGEN (for "Reference Section Generator"), but two are produced by EDGEN and one by ID DICT (for "Identifier Dictionary"). The presentation is organized according to whether the display is textual or tabular. There is also one graphical display. Each document name is underlined on first appearance.

1. Text Documents

The text documents are all word frames (text windows), and will be discussed in this order:

Calling Sequence (REFGEN)
Modular Outline (REFGEN)
Natural Language Summary (REFGEN)
Topological Sort (REFGEN)

The Calling Sequence document has a row for each genus giving the genus name followed by a list of the genera (if any, in order) that it calls directly. (In terms of the genus graph, genus B calls genus A "directly" if there is an arc from A to B, and calls it "indirectly" if there is no arc from A to B and yet A reaches B.) There is also a name-only row for each module to facilitate cross-references with the schema. Figure A.1 gives an example. The information in this document can be obtained from the Adjacency Matrix, which will be discussed later.

&NUT_DATA		
NUTR		
MIN NUTR		
&MATERIALS		
MATERIAL		
UCOST	MATERIAL	
ANALYSIS	NUTR	MATERIAL
Q	MATERIAL	
NLEVEL	ANALYSIS	Q
T:NLEVEL	NLEVEL	MIN
TOTCOST	UCOST	Q

Fig. A.1 Calling Sequence Document for Feedmix Schema

String search provides one of the more useful manipulations of this word frame. For example, it can easily find all direct references to a given genus. Another possibility is to copy this document into a database frame (table window) and then sort on the first column, so that the same information can be accessed alphabetically by genus name. Figure A.2 shows the result, after deleting the two module rows.

GENUS NAME	1st CALL	2nd CALL
ANALYSIS	NUTR	MATERIAL
MATERIAL		
MIN	NUTR	
NLEVEL	ANALYSIS	Q
NUTR		
Q	MATERIAL	
T:NLEVEL	NLEVEL	MIN
TOTCOST	UCOST	Q
UCOST	MATERIAL	

Fig. A.2 Calling Sequence Document for Feedmix Schema After Copying to a Database Frame and Sorting on First Column

The concept of a Modular Outline was introduced in Section 2 of <MS 87> and illustrated in Figures 4, 9, and 14 of that paper. It is automatically available as Framework's "outline view" of the Schema Section, to which there is easy toggle access at all times. This gives a concise representation of the schema's modular structure, and is useful for communication and documentation. Subtree hiding/unhiding, which can be accomplished easily using Framework's outliner, is a particularly useful manipulation. For example, in an on-line briefing one may start with everything hidden except the immediate descendants of the root, and selectively unhide and re-hide detail as appropriate.

The Modular Outline just described is not really a text display. *REFGEN* does, however, produce a true text display: a word frame containing an enhanced Modular Outline that appends the free indices, if any, to each genus name. Figure A.3 shows an example in which the root node has been restored (the root node is implicit in the *REFGEN* document). Compare this with Figure 4 of <MS 87>. One possible manipulation is string search to find all occurrences of a given index.

```

&FEEDMIX
  &NUT DATA
    NUTRi
    MINi
  &MATERIALS
    MATERIALm
    UCOSTm
    ANALYSISim
    Qm
    NLEVELi
    T:NLEVELi
    TOTCOST
  
```

Fig. A.3 Enhanced Modular Outline for Feedmix Schema

The Natural Language Summary was introduced in Section 3.3 of <MS 87> and illustrated in Figure 15 of that paper. As explained there, it is identical to the schema except that it replaces the formal part of each genus paragraph (everything but the interpretation part) by the genus name with free indices appended as in the enhanced Modular Outline produced by *REFGEN*. For example, the genus paragraph

ANALYSIS (NUTRi, MATERIALm) /a/ {NUTR} x {MATERIAL} : Real+ For each NUTRIENT-MATERIAL combination, there is an ANALYSIS (units of nutrient per pound of material).

becomes

ANALYSISim For each NUTRIENT-MATERIAL combination, there is an ANALYSIS (units of nutrient per pound of material).

The Natural Language Summary associated with a schema is especially useful for communication and documentation at a non-mathematical level because it suppresses mathematical details. One possible manipulation would be to tailor it to specific audiences by deleting subtree detail, as in Figure 17 of <MS 87>.

The Topological Sort document contains the so-called "topological labels" discussed at length just after Proposition 7 of Geoffrion <1989b>. As explained in that paper, the topological labels readily identify some, but not necessarily all, of the possible monotone orderings for the schema's modular structure. This could facilitate efforts to find a monotone ordering that is more intuitively pleasing, although the particular ordering used by the current schema is, of course, monotone if the *SMLCheck* process terminated without detecting any errors. An example is shown in Figure A 4. Note that this

<u>SIBLING SET</u>	<u>TOPOLOGICAL LABEL</u>	<u>GENUS RANK</u>
&NUT_DATA	1	0
&MATERIALS	2	0
Q	3	2
NLEVEL	4	3
T:NLEVEL	5	4
TOTCOST	4	3
NUTR	1	1
MIN	2	2
MATERIAL	1	1
UCOST	2	2
ANALYSIS	2	2

Fig. A.4 Topological Sort Document for Feedmix Schema

document also gives the "rank" of each genus, that is, one plus the length of the longest path in the genus graph to the genus in question. This information can be helpful for genus graph layout and as a simple measure of the conceptual depth of the definitions represented by the genera.

Among the useful manipulations of the Topological Sort are the ones that become possible after copying it into a database frame. Then one can easily sort sibling sets individually by topological label, or all genera by rank or alphabetically.

2. Tabular Documents

The tabular documents are as follows, and will be discussed in this order:

Genus/Module Summary (*EDGEN*)
 Identifier Dictionary (*ID_DICT*)
 Matrices (*REFGEN*)
 Adjacency Matrix
 Reachability Matrix
 Adjacency/Reachability Matrix
 Table Structure (*EDGEN*)

The matrices make it very easy to recover the adjacency and reachability matrices for the genus graph and for any module graph (defined in Geoffrion <1989b>). Matrices based on the element graph are not produced.

Figure A.5 gives an example of the Genus/Module Summary. The NAME column contains all genus and module names in modular outline order. The SEQ_NO and PATH columns provide two reference numberings to permit the rows to be restored to their original order after rearrangement and to facilitate finding paragraphs in the schema. The first numbering is strictly sequential. The second one

NAME	SEQ_NO	PATH	TYPE	TABLE	KEY PHRASE
&NUT_DATA	1	1	m		NUTRIENT_DATA
NUTR	2	1.1	pe	NUTR	NUTRIENTS
MIN	3	1.2	a	NUTR	MINIMUM_DAILY_REQUIREMENT
&MATERIALS	4	2	m		MATERIALS_DATA
MATERIAL	5	2.1	pe	MATERIAL	MATERIALS
UCOST	6	2.2	a	MATERIAL	UNIT_COST
ANALYSIS	7	2.3	a	ANALYSIS	ANALYSIS
Q	8	3	va	Q	QUANTITY
NLEVEL	9	4	f	NLEVEL	NUTRITION_LEVEL
T:NLEVEL	10	5	t	NLEVEL	NUTRITION_TEST
TOTCOST	11	6	f	TOTCOST	TOTAL_COST

Fig. A.5 Genus/Module Summary for Feedmix Model

recognizes the sibling sets of modular structure and their ordering (cf. Topological Sort). It encodes, in a standard way, the position of each node in the modular outline. The TYPE column indicates which of the five possible types each genus is, and lists modules as "m". The TABLE column gives the elemental detail table name, if any, corresponding to each genus. Finally, the KEY PHRASE column gives the first underlined key phrase, if any, in the interpretation part of each genus and module paragraph.

Sorting is a very useful manipulation of the Genus/Module Summary. Sorting on NAME produces an alphabetized dictionary of genus and module names, on TYPE produces a list of the modules and a list of the genera grouped by type, on TABLE produces an alphabetized list of tables and the genera associated with each, and on KEY PHRASE produces an alphabetized list of key phrases. Filtering via logical expressions can also be useful.

The Identifier Dictionary gives an exhaustive list of all identifiers appearing in the stubs of elemental detail tables corresponding to self-indexed genera. If the same identifier is introduced by multiple self-indexed genera, then it appears multiple times. Each row gives not only an identifier, but also

the genus which introduced it and its interpretation (if any). For the feedmix model elemental detail tables appearing in Figure 6 of <MS 87>, the result is shown in Figure A.6. One useful manipulation of such a document is to sort it on the "IDENTIFIER" column, so that any duplicates will be evident and alphabetical access will be facilitated (e.g., for checking the meaning of some cryptic identifier).

GENUS	IDENTIFIER	INTERPRETATION
NUTR	P	Protein
NUTR	C	Calcium
MATERIAL	std	Standard Feed
MATERIAL	add	Additive

Fig. A.6 Identifier Dictionary for Feedmix Model

The Adjacency Matrix for the feedmix schema is shown in Figure A.7. The modular outline determines both row and column order. The convention is that "column calls row"; thus, reading by columns indicates which row genera are called by a given column genus, and reading by rows indicates which column genera call a given row genus. These two conventions imply, by the no-forward-reference property of SML schemas, that this matrix always is upper triangular.

Fig. A.7 Adjacency Matrix for Feedmix Schema

Note that modules appear in the matrix as well as genera. That is because FW/SM's Adjacency Matrix is designed to convey not only the adjacency matrix for the genus graph, but also the adjacency matrix for any possible module graph by simply masking out the unwanted rows and columns. We refer the reader to Geoffrion (1989b) for a proper definition of a module graph, and give no further discussion.

Appendix

here concerning the use of the Adjacency Matrix in that connection. An exactly parallel comment applies to FW/SM's Reachability Matrix and Adjacency/Reachability Matrix.

The rules governing the appearance of Adjacency Matrix entries are as follows:

<u>Column</u>	<u>Row</u>	<u>Condition for an Entry at Column-Row</u>
module	module	some genus in the column module directly calls some genus in the row module, and the row module is not identical to or in the column module
module	genus	some genus in the column module directly calls the row genus, and the row genus is not in the column module
genus	module	the column genus directly calls some genus in the row module, and the row module does not contain the column genus
genus	genus	the column genus directly calls the row genus

Fig. A.8 Reachability Matrix for Feedmix Schema

The Reachability Matrix is derived from the Adjacency Matrix in the expected way, with the addition of entries down the diagonal to facilitate reading it. From Figure A.7 one obtains Figure A.8. It follows from the definition of the Adjacency Matrix that the Reachability Matrix obeys these rules for the off-diagonal entries:

<u>Column</u>	<u>Row</u>	<u>Condition for an Entry at Column-Row</u>
module	module	some genus in the column module directly or indirectly calls some genus in the row module, and the row module is not in the column module
module	genus	some genus in the column module directly or indirectly calls the row genus, and the row genus is not in the column module
genus	module	the column genus directly or indirectly calls some genus in the row module, and the row module does not contain the column genus
genus	genus	the column genus directly or indirectly calls the row genus

Reading the Reachability Matrix by columns indicates which genera and modules definitionally influence the column genus or module. Reading by rows indicates which genera and modules are influenced definitionally by the row genus or module. Thus, for example, if an error is discovered in an attribute genus' element value, one can tell at a glance which other function or test genus values could possibly be affected.

The Adjacency/Reachability Matrix is a version of the Reachability Matrix in which the entries are "0" on the diagonal, "1" if the corresponding call is direct, and "2" if the corresponding call is indirect. Thus, one may easily reconstruct both the Adjacency Matrix and the Reachability Matrix from this one matrix. Figure A.9 shows an example.

	&							
	M							
N	A	M	A					T
U	T	A	N				:	T
T	E	T	A			N	N	O
	R	E	U	L		L	L	T
D	N	I	R	C	Y	E	E	C
A	U	M	A	I	O	S	V	V
T	T	I	L	A	S	I	E	S
A	R	N	S	L	T	S	Q	L
								T
&NUT_DATA								
NUTR	0	.	.	1	.	.	1	.
MIN	.	0	1
&MATERIALS								
MATERIAL	.	.	.	0	.	.	1	1
UCOST	0	.	.	1
ANALYSIS	0	.	1
Q	0	1	2
NLEVEL	0	1
T:NLEVEL	0
TOTCOST	0

Fig. A.9 Adjacency/Reachability Matrix for Feedmix Schema

Before discussing a few of the easy but useful manipulations of these matrices, we must note that *REFGEN* actually produces them as word frames. However, the DBREF process can automatically convert

them to database frames. Thus one may use either Framework's word processor alone on a word frame or together with Framework's table editor on a database frame.

Sorting the rows on the row name column of the database frame version of a matrix puts the rows in alphabetical order according to row name. This facilitates row access by genus or module name. Sorting the rows on any other column groups rows with identical entries in the sort column; in this way one can, for example, easily produce a list of all genera and modules that reach a given genus or module. Filtering an Adjacency/Reachability Matrix can be used to identify genera and modules which satisfy any desired logical criteria involving their direct or indirect participation in specified genera and modules; for example, one can easily identify all genera which are called directly or indirectly by any of a list of specific modules.

The last tabular document is Table Structure. It gives the names of all elemental detail tables, in order, and their columns. It also gives the "type" of each table, a technical attribute defined in Geofrion <1988> that ordinary users can ignore. Figure A.10 gives an example; examples of all of these tables can

NAME	STRUCTURE	TYPE
NUTR	NUTR INTERP MIN	1
MATERIAL	MATERIAL INTERP UCOST	1
ANALYSIS	NUTR MATERIAL ANALYSIS	2
Q	MATERIAL Q	2
NLEVEL	NUTR NLEVEL T:NLEVEL	2
TOTCOST	TOTCOST	4

Fig. A.10 Table Structure Document for Feedmix Model

be found in Figure 6 of <MS 87>, and an example of the third table appears at the end of Section 2.3 of the present paper. Sorting the rows on the column "NAME" produces an alphabetized list of tables and their columnar structure. Filtering can find the tables with specific characteristics; for example, it can find all tables containing a specific genus name (use Framework's @scan function).

3. Graphical Display

The main graphical display of interest is the Genus Graph. Figures 3, 8, and 12 of <MS 87> and Section 3.3 thereof give examples and explain its importance for communication purposes. The NETDRAW process is an interactive implementation in the style of Javelin's "diagram view" (Javelin <1986>) that allows any one node of the genus graph to be displayed centrally on screen along with all adjacent arcs and nodes. The genus graph can be "walked" with the cursor keys, which move a highlight that can change the central node to one of its neighbors.

The modular tree (e.g., Figure 4 of <MS 87>) is not implemented, although the Modular Outline is equivalent and available in two forms as mentioned previously.